Config2Spec: Mining Network Specifications from Network Configurations



Rüdiger Birkner, Dana Drachsler-Cohen, Martin Vechev, Laurent Vanbever

nsg.ee.ethz.ch

NSDI'20 February, 27 2020



ETHzürich

Many tools are available that allow you to check that your network behaves as intended



Standard recipe:

- **Upload configurations**
- Define specification 2
- Run the tool 3
- Iterate & deploy



Many tools are available that allow you to check that your network behaves as intended



Standard recipe:

- **Upload configurations**
- Define specification 2
- Run the tool 3
- Iterate & deploy



Definition The specification of a **network** is the **set of all policies** that hold...

Set of policies
 reachability(r1,p1)
 waypoint(r3,r1,p2)
 reachability(r5,p2)

loadbalancing(r3,p2)







Definition The specification of a **network** is the **set of all policies** that hold...

Set of policies





reachability(r5,p2)

loadbalancing(r3,p2)







Definition The specification of a **network** is the **set of all policies** that hold...

Set of policies



reachability(rs,p2)

loadbalancing(r3,p2)





Definition The specification of a **network** is the **set of all policies** that hold under a given **failure model**.

Set of policies
 reachability(r1,p1)
 waypoint(r3,r1,p2)
 reachability(r5,p2)

loadbalancing(r3,p2)



Definition The specification of a **network** is the set of all policies that hold under a given failure model.



Set of concrete environments



Symbolic environment

Failure bound

k = 2



Definition The specification of a **network** is the **set of all policies** that hold under a given **failure model**.

Set of policies Sy
reachability(r1,p1)
waypoint(r3,r1,p2)
reachability(r5,p2)

loadbalancing(r3,p2)

Symbolic environment

Failure bound



k = 2



Writing the network's precise specification is hard



```
لاجي خلاف) Dr Heidy Khlaaf.
```

@HeidyKhlaaf

In the past three years of working on large safety critical systems, I've learned that verification isn't the real problem, but it's writing specifications. Don't @ me. twitter.com/Conaw/status/1...

Putting network verification to good use

Ryan Beckett Microsoft Research

Ratul Mahajan University of Washington Intentionet

... However, outside of a handful of large cloud computing providers, the use of network verification is still sparse.



Standard recipe:

- Upload configurations
- 2 Define specification
- Run the tool 3
- Iterate & deploy

Internet2's specification with its 10 routers consists of ~4000 policy predicates.

Internet2's specification with its 10 routers consists of ~4000 policy predicates.

Imagine writing that specification by hand.

Introducing

Config2Spec



Config2Spec automatically mines the network's full specification from its configuration and the given failure model



Config2Spec

Output

Network Specification

loadbalancing(4, p2)
reachability(1, p1)
reachability(1, p2)

...

reachability(4, p2)
reachability(5, p2)

Config2Spec: Mining Network Specifications from Network Configurations

3

1

Baseline approaches one search space at a time

Our approach the best of both worlds

Evaluation scales to realistic networks

Config2Spec: Mining Network Specifications from Network Configurations

3

1

Baseline approaches one search space at a time

Our approach the best of both worlds

Evaluation scales to realistic networks



Mining a network specification involves exploring two exponential search spaces



all concrete environments

reachability(r1,p1)

waypoint(r3,r1,p2)



all possible policies

Mining a network specification involves exploring two exponential search spaces

data plane analysis

control plane verification



data plane analysis





Data plane analysis tools allow to find all the policies that hold for a single concrete environment



Concrete environment

Data plane analysis



The network specification is the intersection of the policies that hold for every concrete environment



. . .

. . . .

21

The network specification is the intersection of the policies that hold for every concrete environment





data plane analysis



23

Control plane verification tools determine whether a policy holds for the entire failure model



Control plane verification

Verification result

24

The network specification is the set of policies that the verifier determined to hold for the failure model



p1) •••

...

25

The network specification is the set of policies that the verifier determined to hold for the failure model



under-approximation of the specification





Both techniques have pros and cons

approach data plane and all policies for one concrete

good at pruning dense violation

nalysis	control plane verification
or e env.	one policy for the entire failure model
ons	sparse violations

27

Violations are policies that are not part of the specification





- dense violation
- no envs.
- waypoint(r5,r2,p2)

sparse violation all envs.

reachability(r1,p2)







28

Both techniques have pros and cons

approach data plane ar all policies fo one concrete

good at pruning dense violations sparse violations

nalysis	control plane verification
or e env.	one policy for the entire failure model

What about combining them?

29

Config2Spec: Mining Network Specifications from Network Configurations

3

1

Baseline approaches one search space at a time

Our approach the best of both worlds

Evaluation scales to realistic networks



Config2Spec mines the network's full specification from its configuration and the required failure tolerance





Thanks to combining the two approaches, Config2Spec is precise

Data plane analysis

over-approximation



Control plane verification under-approximation



32

Step-by-step from **all** existing policies to the network's specification



By performing data plane analysis on a topology, Config2Spec refines the space of candidate policies





By performing data plane analysis on a topology, Config2Spec refines the space of candidate policies





By performing data plane analysis on a topology, Config2Spec refines the space of candidate policies











































When Config2Spec terminates, it is left with the specification.





Config2Spec can be improved further by two domain-specific techniques

policy-aware selection

grouping and trimming



policy-aware selection

grouping and trimming



Data plane analysis has to reduce the candidate set to a minimum as fast as possible

basic	rando
policy-aware	pick

Candidate set
reachability(r1,p2)
reachability(r2,p2)

env #n

Failure model





omly pick concrete environments

concrete envs. based on the candidate set





policy-aware selection

grouping and trimming



To be fast, control plane verification should be used as little as possible



leverage topology and failure model can connectivity requirements be met?

grouping

trimming

verify similar policies at once grouping is based on the destination

verify each policy separately



Config2Spec: Mining Network Specifications from Network Configurations

3

1

Baseline approaches one search space at a time

Our approach the best of both worlds

Evaluation scales to realistic networks



Question #1

How do the domain-specific techniques Question #2 contribute to Config2Spec?

How does Config2Spec scale to large-scale configurations?

53

We fully implement Config2Spec and show its practicality

Implementation

Methodology

- 5k lines of Python and Java relying on Batfish [NSDI'15] and Minesweeper [Sigcomm'17]
- generated configs using NetComplete [NSDI'18] employing OSPF, BGP
- for a small, medium, and large network with 33, 70, and 158 routers





Config2Spec mines the specification for realistic networks in few hours







For failure models with few concrete environments, data plane analysis on its own provides fastest progress







For failure models with a high failure bound, policy trimming reduces the candidate space significantly





Config2Spec mines the specification for realistic networks in few hours



58

Policy-aware environment selection leads to smaller candidate sets with fewer samples than random



of environments



Policy trimming and grouping allows to significantly reduce the number of queries

number of queries for control plane verification

after data plane analysis

after policy trimming

after policy grouping - - - - -



medium topology, k = 2



Config2Spec: Mining Network Specifications from Network Configurations

Baseline approaches one search space at a time

Our approach the best of both worlds

Evaluation scales to realistic networks



Config2Spec: Mining Network Specifications from Network Configurations



automatically learns a network's specification based on its configuration and failure model

scales thanks to the combination of the two approaches data plane analysis and control plane verification

the specification is useful beyond verification what-if analysis, config streamlining, network understanding





Config2Spec Mining Network Specifications from Network Configurations





Rüdiger Birkner

Dana Drachsler-Cohen



Martin Vechev



Laurent Vanbever

nsg.ee.ethz.ch